

On the Configuration-LP of the Restricted Assignment Problem*

Klaus Jansen

Lars Rohwedder

Department of Computer Science, University of Kiel, 24118 Kiel, Germany
 {kj, lro}@informatik.uni-kiel.de

Abstract

We consider the classical problem of SCHEDULING ON UNRELATED MACHINES. In this problem a set of jobs is to be distributed among a set of machines and the maximum load (makespan) is to be minimized. The processing time p_{ij} of a job j depends on the machine i it is assigned to. Lenstra, Shmoys and Tardos gave a polynomial time 2-approximation for this problem [8]. In this paper we focus on a prominent special case, the RESTRICTED ASSIGNMENT problem, in which $p_{ij} \in \{p_j, \infty\}$. The configuration-LP is a linear programming relaxation for the RESTRICTED ASSIGNMENT problem. It was shown by Svensson that the multiplicative gap between integral and fractional solution, the integrality gap, is at most $2 - 1/17 \approx 1.9412$ [11]. In this paper we significantly simplify his proof and achieve a bound of $2 - 1/6 \approx 1.8333$. As a direct consequence this provides a polynomial $(2 - 1/6 + \epsilon)$ -estimation algorithm for the RESTRICTED ASSIGNMENT problem by approximating the configuration-LP. The best lower bound known for the integrality gap is 1.5 and no estimation algorithm with a guarantee better than 1.5 exists unless $P = NP$.

Keywords: estimation algorithms, scheduling, local search, linear programming, integrality gap

1 Introduction

In the problem SCHEDULING ON UNRELATED MACHINES we are given a set of jobs \mathcal{J} and a set of machines \mathcal{M} . Each job has to be assigned to exactly one machine. A solution is a function $\sigma : \mathcal{J} \rightarrow \mathcal{M}$. Every Job $j \in \mathcal{J}$ has a processing time p_{ij} , that depends on the machine $i \in \mathcal{M}$ it is assigned to. The objective is to minimize the highest load among all machines, that is $\max_{i \in \mathcal{M}} \sum_{j \in \sigma^{-1}(i)} p_{ij}$. This value is called the makespan of σ . The best polynomial algorithm known for SCHEDULING ON UNRELATED MACHINE has an approximation guarantee of 2 [8]. In the same paper it was also shown that approximating the problem with a rate better than 1.5 is NP-hard. Williamson and Shmoys [13] as well as Schuurman and Woeginger [10] asked as an important open question, whether this can be improved. In both cases, the RESTRICTED ASSIGNMENT problem is mentioned as well. The RESTRICTED ASSIGNMENT problem is a natural special case. There we have $p_{ij} \in \{p_j, \infty\}$ for a processing time p_j which depends only on the job. The intuition is that each job is allowed only on a subset of machines $\Gamma(j) = \{i \in \mathcal{M} : p_{ij} < \infty\}$, but the processing time on any of these machines is identical. The known bounds on approximability of the RESTRICTED ASSIGNMENT match the general problem.

An important instrument for the RESTRICTED ASSIGNMENT problem is a linear programming formulation called the configuration-LP. Svensson has shown that this relaxation has an integrality gap of no more than $2 - 1/17$ [11]. Since the linear program can be solved up to an error of an arbitrarily small ϵ in polynomial time, this gives us a polynomial $(2 - 1/17 + \epsilon)$ -estimation algorithm. A c -estimation algorithm stands for an algorithm that computes a value E with $\text{OPT} \leq E \leq c \cdot \text{OPT}$, where OPT denotes the optimum. In contrast, a c -approximation algorithm would also need to produce a schedule that has a makespan in this range. It is already NP-hard to estimate the RESTRICTED ASSIGNMENT problem with a rate better than 1.5 (as can be shown using the reduction from [8]).

*Research was supported by German Research Foundation (DFG) project JA 612/15-1

Further progress has been made on the special case of two different processing times, where an upper bound of $2 - 1/3$ was shown for the integrality gap by Land, Maack, and the first author [7] and a $(2 - \delta)$ -approximation for a very small δ by Chakrabarty, Khanna, and Li [3].

The method used to show the bound on the integrality gap is a local search algorithm that produces a solution of the mentioned quality. It is, however, not known, whether the algorithm terminates in polynomial time. An indication of the potential of this research is given by the closely related RESTRICTED MAX-MIN FAIR ALLOCATION problem. This problem maximizes the minimum load instead of minimizing the maximum. Similar techniques have been employed for this problem and the respective local search algorithm has been improved first to a quasi-polynomial running time by Poláček and Svensson [9] and eventually to a polynomial one by Annamalai, Kalaitzis, and Svensson [1].

Another noteworthy branch of research regarding the RESTRICTED ASSIGNMENT problem is the special case of GRAPH BALANCING. In this problem each job is allowed on at most two machines, that is for all $j \in \mathcal{J}$ we have $|\Gamma(j)| \leq 2$. Using linear programming techniques a 1.75-approximation was achieved by Ebenlendr, Krcál, and Sgall [4]. Those techniques, however, differ very much from those related to the configuration-LP. Recently Huang and Ott were able to show some further results on instances of RESTRICTED ASSIGNMENT where the constraint above applies only to big jobs [6].

Our contribution. As a direction for future research Svensson stated in his paper [11]:

“[To improve the upper bound on the integrality gap] one possibility would be to find a more elegant generalization of the techniques, presented [...] for two job sizes, to arbitrary processing times (instead of the exhaustive case distinction presented in this paper).”

This accurately captures where our contribution lies. We improve the upper bound on the integrality gap by coping with arbitrary many job sizes in a different way. We also argue that our proof is simpler than the original one (see Section 1.1).

Theorem 1.1. *The configuration-LP for the RESTRICTED ASSIGNMENT problem has an integrality gap of at most $2 - 1/6$.*

Corollary 1.2. *For every $\epsilon > 0$ there exists a polynomial $(2 - 1/6 + \epsilon)$ -estimation algorithm for the RESTRICTED ASSIGNMENT problem.*

1.1 Differences to Svensson’s algorithm

We distinguish the jobs sizes only into small and big, whereas the original proof additionally has a fine distinction of big jobs into medium, large, and huge.

One of the difficulties that arise in Svensson’s paper when dealing with many job sizes is the role of the medium jobs. There are situations where we want to move huge jobs to certain machines, but these machines contain too many medium jobs. In the process of removing those medium jobs, it is hard to prevent bigger medium jobs from being moved to this machine. The original paper solves this by rounding the medium jobs to the same size. To be precise, the rounding affects only the values in the construction of a dual LP solution, which is used in the analysis. In contrast, we accept that during this process new medium jobs are moved to the machine, as long as they are bigger than the smallest one on the machine. Our observation is that when repeating this process the size of the smallest medium job steadily increases; thus the process eventually terminates.

Rounding of medium jobs introduces some unpleasant problems. Our approach avoids these problems making the analysis a lot less technical. It also appears to be a bottleneck for the solution quality in the original local search algorithm and with our methods the approximation ratio can be improved significantly.

1.2 Notation

For jobs we usually use the variable j or variants like j_B, j' , etc. For machines we use i and the like. Since we frequently sum over sets of jobs, for a $H \subseteq \mathcal{J}$ we write $\sum_{j \in H} p_j$ simply as $p(H)$. With other variables indexed by jobs we do the same.

$ \begin{aligned} \sum_{C \in \mathcal{C}_i(T)} x_{i,C} &\leq 1 & \forall i \in \mathcal{M} \\ \sum_{i \in \mathcal{M}} \sum_{C \in \mathcal{C}_i(T): j \in C} x_{i,C} &\geq 1 & \forall j \in \mathcal{J} \\ x_{i,C} &\geq 0 \end{aligned} $	$ \begin{aligned} \min \quad & \sum_{i \in \mathcal{M}} y_i - \sum_{j \in \mathcal{J}} z_j \\ s.t. \quad & \\ & y_i \geq \sum_{j \in C} z_j \quad \forall i \in \mathcal{M}, C \in \mathcal{C}_i(T) \\ & y_i, z_j \geq 0 \end{aligned} $
(a) Primal	(a) Dual

Figure 3: The configuration-LP

1.3 The configuration-LP

A configuration for a machine is a set of jobs that would not exceed the target makespan T , if the jobs were assigned to that machine. The set of all configurations for a machine i is therefore defined as

$$\mathcal{C}_i(T) = \{C \subseteq \mathcal{J} : \sum_{j \in C} p_{ij} \leq T\}.$$

Note that by this definition a job j will only appear in configurations of machines in $\Gamma(j)$. The idea for the configuration-LP (Figure 3) is that for each machine a mixture of configurations is selected. A variable $x_{i,C}$ describes which fraction of the configuration C is used on machine i . The first constraint guarantees that the sum of these fractions on each machine is at most 1. The second constraint ensures that each job appears in fractional configurations with a total value of at least 1. It is easy to see that an integral solution to the configuration-LP directly correlates with a solution for the respective instance of the RESTRICTED ASSIGNMENT problem, when T equals the optimum. The minimal T for which the LP is feasible is considered the optimum of the relaxation and is denoted by OPT^* . Although the configuration-LP has an exponential number of variables, it can be approximated in polynomial time to a factor of $1 + \epsilon$ for any $\epsilon > 0$ [2]. A well-known result from linear programming theory is that an unbounded dual implies an infeasible primal. This will be a central argument in our analysis.

We note that in the more general SCHEDULING ON UNRELATED MACHINE problem the configuration-LP has an integrality gap that cannot be bounded by a constant lower than 2 [12], which indicates that these techniques are probably not particularly helpful there. In the restricted variant, however, no instance is known to have an integrality gap higher than 1.5. An instance with integrality gap 1.5 is given in [7].

1.4 Reduced problem

Let us first simplify our problem. The makespan we aim for is $(1 + R)\text{OPT}^*$, where $R = 5/6$. For convenience we scale each processing time by $1/\text{OPT}^*$; thereby we establish $\text{OPT}^* = 1$. Unless otherwise stated, statements about configurations or the configuration-LP are meant in respect to a makespan of 1.

Definition 1.3 (Small and big jobs). A job j is *small* if $p_j \leq 1/2$ and *big* otherwise. The set of small (big) jobs is denoted by \mathcal{J}_S (respectively, \mathcal{J}_B).

A crucial difference between small and big jobs is that there can be more than one small job in a configuration, whereas at most one big job may occur.

Definition 1.4 ([Valid] partial schedule). We call a function $\sigma : \mathcal{J} \rightarrow \mathcal{M} \cup \{\perp\}$ a partial schedule when for each job j either $\sigma(j) = \perp$ or $\sigma(j) \in \Gamma(j)$. The value \perp is used to indicate that a job has not been assigned, yet. A partial schedule is said to be *valid* if for each machine $i \in \mathcal{M}$ the sum of the processing times of the jobs assigned to it is at most $1 + R$, that is $p(\sigma^{-1}(i)) \leq 1 + R$.

To find a schedule for all jobs, we maintain a valid partial schedule and extend it one job at a time (in an arbitrary order). The problem that we will focus on in this paper is given below.

Input: An instance of RESTRICTED ASSIGNMENT with $\text{OPT}^* = 1$, a valid partial schedule σ , and a job j_{new} with $\sigma(j_{\text{new}}) = \perp$.

Output: A valid partial schedule σ' with $\sigma'(j_{\text{new}}) \neq \perp$ and if $\sigma(j) \neq \perp$ then $\sigma'(j) \neq \perp$ for all $j \in \mathcal{J}$.

Without loss of generality assume that the job identifiers are natural numbers, i.e., $\mathcal{J} = \{1, 2, 3, \dots\}$, and the jobs are ordered by size ($p_1 \leq p_2 \leq \dots$). This provides a total order on the jobs and simplifies the notation in algorithm and analysis.

2 Local search algorithm

The key idea for the algorithm is a structure of so-called blockers. We motivate it with the thought below. Suppose adding j_{new} exceeds the capacity of every machine in $\Gamma(j_{\text{new}})$. This means we have to move away some jobs from those machines first. Indeed, these jobs now have a similar role to j_{new} . In other words, we try to find a new machine for them as well. We can repeat this with the hope that at some point we can move one of the jobs. To keep track of all dependencies between machines and jobs we use blockers, which represent moves that are placed on hold to be performed at a later time.

Definition 2.1 ([Valid] moves). A move (j, i) (with respect to a partial schedule σ) is composed of a job j and a machine i requiring that j is allowed on i , but is assigned to a different machine, that is $i \in \Gamma(j) \setminus \{\sigma(j)\}$. We call a move (j, i) *valid* if $p(\sigma^{-1}(i)) + p_j \leq 1 + R$.

Note that for a valid move (j, i) altering the assignment of j to i does not compromise the validity of the partial schedule. Here and subsequently, B_i denotes the set of big jobs on machine i , i.e., $\sigma^{-1}(i) \cap \mathcal{J}_B$. We define

$$B_i^{\min} := \begin{cases} \{\min B_i\} & \text{if } B_i \neq \emptyset \text{ and} \\ \emptyset & \text{otherwise.} \end{cases}$$

Recall that jobs are represented by natural numbers and are ordered by size. As opposed to the rather natural set $\{j \in B_i \mid \forall j' \in B_i (p_j \leq p_{j'})\}$, B_i^{\min} always has exactly one or zero elements. Also, B_i^{\min} is always defined, whereas $\min B_i$ is undefined for $B_i = \emptyset$. Note that S_i , which we will introduce later as well, is not analogous to B_i .

Definition 2.2 (Blockers). A blocker is a tuple (j, i, Θ) where (j, i) is a move and Θ is the type of the blocker. There are four types with the following abbreviations:

(SA) *small-to-any blockers*, (BA) *big-to-any blockers*, (BB) *big-to-big blockers*, (BL) *big-to-least blockers*.

As the first part of the name suggests, each small-to-any blocker (j, i, SA) corresponds to a small job j ; all other types correspond to big jobs. The latter part of the type's name is a hint at which jobs are marked undesirable on the machine corresponding to the blocker. The algorithm will try to move jobs that are undesirable on a machine away from it; at the same time it will not attempt to move such jobs onto the machine. On machines of small-/big-to-any blockers all jobs are undesirable; on machines of big-to-big blockers big jobs are undesirable; finally on machines of big-to-least blockers only particular big jobs are undesirable. The goal for this type is to steadily increase the size of the smallest big job. To achieve this we would like to remove the smallest job from B_i and at the same time ensure that no smaller jobs are being added to B_i (big-to-least blockers will ensure that $B_i \neq \emptyset$). More precisely, all big jobs of index smaller or equal to $\min B_i$ are marked undesirable.

The blocker tree \mathcal{T} is a set of blockers. We will elaborate the tree analogy in a moment. Roughly speaking, the blockers correspond to moves that the algorithm considers useful for assigning j_{new} . For simplicity of notation, we say a move (j, i) is in \mathcal{T} , if there is a blocker (j, i, Θ) for some type Θ with $(j, i, \Theta) \in \mathcal{T}$. The blockers corresponding to the specific types are written as \mathcal{T}_{SA} , \mathcal{T}_{BA} , \mathcal{T}_{BB} , and \mathcal{T}_{BL} .

With \mathcal{T} we associate a tree structure, in which the blockers and one additional root form the nodes. Each blocker $\mathcal{B} = (j, i, \Theta)$ has a parent, that is determined by j . It is a blocker for machine $\sigma(j)$ with a type, for which j is regarded undesirable, or the root if $j = j_{\text{new}}$. If there are several candidates for a parent, the algorithm uses the one that was added to the blocker tree first. We say this blocker activates the job (see also definition of active jobs below). It is important to understand that the tree is always connected,

Algorithm 1: Local search

```

// Input: Job  $j_{\text{new}}$  and partial schedule  $\sigma$ 
initialize empty blocker tree  $\mathcal{T}$ ;
loop
  if a move  $(j, i)$  in  $\mathcal{T}$  is valid then
    Let  $\mathcal{B}_k$  be the blocker that activated  $j$ ;
    // Update the schedule
     $\sigma(j) \leftarrow i$ ;
    if  $j = j_{\text{new}}$  then
      return  $\sigma$ ;
    end
    // Delete  $\mathcal{B}_k, \mathcal{B}_{k+1}, \dots$ 
     $\mathcal{T} \leftarrow \mathcal{T}^{(\leq k-1)}$ ;
  else
    choose a potential move  $(j, i)$  with minimum value;
    add  $(j, i)$  to  $\mathcal{T}$  with the correct type;
  end
end

```

i.e., a parent exists for every blocker. This, however, becomes obvious once we complete the description the algorithm.

We denote the individual blockers by $\mathcal{B}_1, \dots, \mathcal{B}_\ell$ in the order they were added. For $k \leq \ell$ we write $\mathcal{T}^{(\leq k)}$ for the blockers $\mathcal{B}_1, \dots, \mathcal{B}_k$. In the final algorithm whenever we remove a blocker \mathcal{B}_k , we also remove all blockers added to the tree after it, that is $\mathcal{B}_{k+1}, \dots, \mathcal{B}_\ell$.

From the blocker tree, we derive the machine set $\mathcal{M}(\mathcal{T})$ which consists of all machines corresponding to blockers in \mathcal{T} . The same notation is used for subsets of \mathcal{T} , such as for instance in $\mathcal{M}(\mathcal{T}^{(\leq k)})$ and $\mathcal{M}(\mathcal{T}_{\text{BA}})$.

The same machine can appear more than once in the blocker tree. In that case, the undesirable jobs are the union of the undesirable jobs from all types. Also, the same job can appear multiple times in different blockers.

Definition 2.3 (Active jobs). We consider the small jobs j that are undesirable on all other machines they are allowed on, that is

$$\Gamma(j) \setminus \{\sigma(j)\} \subseteq \mathcal{M}(\mathcal{T}_{\text{SA}} \cup \mathcal{T}_{\text{BA}}).$$

We write $S(\mathcal{T})$ for these jobs. They are small jobs that we cannot hope to move at the moment. By $S_i(\mathcal{T})$ we denote for a machines i the set of jobs in $S(\mathcal{T})$ that are assigned to i , that is $S_i(\mathcal{T}) := S(\mathcal{T}) \cap \sigma^{-1}(i)$. The set of *active* jobs $\mathcal{A}(\mathcal{T})$ includes j_{new} , $S(\mathcal{T})$ as well as all those jobs, that are undesirable on the machine they are currently assigned to. Like with $S_i(\mathcal{T})$, we denote by $\mathcal{A}_i(\mathcal{T})$ the set $\mathcal{A}(\mathcal{T}) \cap \sigma^{-1}(i)$.

2.1 Detailed description of the algorithm

The general outline of the algorithm is that as long as j_{new} has not been assigned, the algorithm (see Algorithm 1) performs a valid move in the blocker tree, if possible, and otherwise adds a new blocker.

Adding blockers. The move, that is added to \mathcal{T} , has to meet certain requirements. A move, that does, is called a potential move. For each type of blocker we define a type of potential move: Potential big-to-any moves, potential big-to-big moves, etc. When a potential move is added to the blocker tree, its type will then be used for the blocker. Let $j \in \mathcal{A}(\mathcal{T})$. For a move (j, i) to be a potential move of a certain type, it has to meet the following requirements.

1. (j, i) is not already in \mathcal{T} .
2. The size of j corresponds to the type, i.e., if j is small, (j, i) must be a small-to-any move and if j is big, (j, i) must be a big-to-any, big-to-big, or big-to-least move.

Table 1: Types of potential moves / blockers w.r.t. a move (j, i)

Type	Conditions	$\text{val}(j, i)$	Undesirable
Small-to-any (SA)	None	$(1, \sigma^{-1}(i))$	All jobs
Big-to-any (BA)	$p(S_i(\mathcal{T}) \cup B_i) + p_j \leq 1 + R$	$(2, \sigma^{-1}(i))$	All jobs
Big-to-least (BL)	$p(S_i(\mathcal{T}) \cup B_i^{\min}) + p_j > 1 + R$	$(3, -\min B_i)$	Big jobs j_B with $j_B \leq \min B_i$
Big-to-big (BB)	$p(S_i(\mathcal{T}) \cup B_i) + p_j > 1 + R$ $p(S_i(\mathcal{T}) \cup B_i^{\min}) + p_j \leq 1 + R$	$(4, B_i)$	Big jobs


3. j is not undesirable on i . In other words, $i \notin \mathcal{M}(\mathcal{T}_{\text{SA}} \cup \mathcal{T}_{\text{BA}})$ and if j is big, then $i \notin \mathcal{M}(\mathcal{T}_{\text{BB}})$ and either $i \notin \mathcal{M}(\mathcal{T}_{\text{BL}})$ or $\min B_i < j$.
4. The load on i meets certain conditions depending on the type. For the sake of readability, these conditions are arranged in table form (see Table 1).

If we compare the conditions in the table, we notice that for moves of big jobs, there is exactly one type that applies, unless $p(S_i(\mathcal{T})) + p_j > 1 + R$, in which case the move is never a potential move. The table also lists a value for each type of move. The algorithm will add the move that has the lowest value (in lexicographical order). It may be useful at this point to take a closer look at the value of a big-to-least move, that is $(3, -\min B_i)$. Here $\min B_i$ is a certain job. The definition only makes sense, because we use the convention that jobs are identified by natural number $\{1, \dots, |\mathcal{J}|\}$.


Performing valid moves. If a move (j, i) in \mathcal{T} is valid, it can be performed. Let \mathcal{B}_k be the blocker that activated j . The algorithm will simply reassign j to i and discard $\mathcal{B}_k, \mathcal{B}_{k+1}, \dots$. The idea is that during the next iteration of the loop, \mathcal{B}_k could be added back to \mathcal{T} , potentially with a different type.

2.2 Example

This example demonstrates the process of a move being performed. In the illustrations the dashed boxes show the capacity of a machine, hatched jobs are those, which the blocker activates, and the gray job is the



(a) The blocker tree before the move,



(b) after the move, and (c) after the blocker is added back.

one that will be moved. The initial state is given in (a). In the current blocker tree 3 different machines are involved. The machine corresponding to \mathcal{B}_1 is the same as the one for \mathcal{B}_4 . It was first added as a big-to-least blocker and then later as a small-to-any blocker. The move for \mathcal{B}_4 is valid. The job was activated by \mathcal{B}_2 . After the job is moved (see (b)), the blockers starting with the one that activated the gray job are removed from the blocker tree. Since the move for \mathcal{B}_2 is still a potential move, it can be added back to the blocker tree in the next iteration (see (c)). We notice that the value of the move, that is $(2, |\sigma^{-1}(M_2)|)$, has decreased between (a) and (c).

3 Analysis

The correctness of the algorithm relies on two theorems. First, we need to understand that at any point the blocker tree either contains a valid move, which can be performed, or there is a potential move that can be added to the blocker tree. This means that the algorithm does not get stuck at any point during its execution. Next, we show that after a bounded number of steps the algorithm terminates.

Lemma 3.1 (Invariants). *At the start of each iteration of the loop,*

1. for all $(j, i, \text{BA}) = \mathcal{B}_{k+1} \in \mathcal{T}_{\text{BA}}: p(S_i(\mathcal{T}^{(\leq k)}) \cup B_i) + p_j \leq 1 + R;$
2. for all $(j, i, \text{BL}) = \mathcal{B}_{k+1} \in \mathcal{T}_{\text{BL}}: p(S_i(\mathcal{T}^{(\leq k)}) \cup B_i^{\min}) + p_j > 1 + R$ and $p(S_i(\mathcal{T}^{(\leq k)})) + p_j \leq 1 + R;$
3. for all $(j, i, \text{BB}) = \mathcal{B}_{k+1} \in \mathcal{T}_{\text{BB}}: p(S_i(\mathcal{T}^{(\leq k)}) \cup B_i) + p_j > 1 + R$ and $p(S_i(\mathcal{T}^{(\leq k)}) \cup B_i^{\min}) + p_j \leq 1 + R.$

Proof. During the lifespan of a blocker $\mathcal{B}_{k+1} = (j, i, \Theta)$, $\mathcal{T}^{(\leq k)}$ does not change; otherwise \mathcal{B}_{k+1} would be deleted. By definition, we cannot add potential moves for jobs in $S_i(\mathcal{T}^{(\leq k)})$ as long as $\mathcal{T}^{(\leq k)}$ does not change. At the same time no blockers for small jobs in $\sigma^{-1}(i)$ can already be in $\mathcal{T}^{(\leq k)}$, since otherwise $i \in \mathcal{M}(\mathcal{T}_{\text{SA}}^{(\leq k)} \cup \mathcal{T}_{\text{BA}}^{(\leq k)})$ and \mathcal{B}_{k+1} could not have been added. For this reason we can deduce that $p(S_i(\mathcal{T}^{(\leq k)}))$ is constant during the lifespan of \mathcal{B}_{k+1} . When a job, which was activated by \mathcal{B}_{k+1} or an earlier blocker, is removed from i , then \mathcal{B}_{k+1} is deleted. For big-to-any or big-to-big blockers, all jobs in B_i are activated within $\mathcal{T}^{(\leq k+1)}$ and therefore do not change during the lifespan of \mathcal{B}_{k+1} . Likewise, for big-to-least blockers, B_i^{\min} does not change. We conclude that the invariants are correct. \square

Theorem 3.2. *At the start of each iteration of the loop, there is a valid move in the blocker tree or a potential move of a job in $\mathcal{A}(\mathcal{T})$.*

Proof. Suppose toward contradiction no potential move of a job in $\mathcal{A}(\mathcal{T})$ remains and no move in the blocker tree is valid. We will show that the dual of the configuration-LP is unbounded (regarding makespan 1), which implies the primal is infeasible. For that purpose, define the solution z^*, y^* below for the dual. For all jobs j and machines i let

$$z_j^* = \begin{cases} \min\{p_j, \frac{5}{6}\} & \text{if } j \in \mathcal{A}(\mathcal{T}), \\ 0 & \text{otherwise;} \end{cases} \quad y_i^* = \begin{cases} z^*(\mathcal{A}_i(\mathcal{T})) + \frac{1}{6} & \text{if } i \in \mathcal{M}(\mathcal{T}_{\text{BA}}), \\ z^*(\mathcal{A}_i(\mathcal{T})) - \frac{1}{6} & \text{if } i \in \mathcal{M}(\mathcal{T}_{\text{SA}}), \\ z^*(\mathcal{A}_i(\mathcal{T})) & \text{otherwise.} \end{cases}$$

Note that since $\mathcal{M}(\mathcal{T}_{\text{BA}}) \cap \mathcal{M}(\mathcal{T}_{\text{SA}}) = \emptyset$, y^* is well-defined: On a blocker of either type, all jobs are undesirable, that is to say as long as one of such blockers remains in the blocker tree, the algorithm will not add another blocker with the same machine.

Claim 3.2.1. *The objective value of the solution is negative, that is $\sum_{j \in \mathcal{J}} z_j^* > \sum_{i \in \mathcal{M}} y_i^*$.*

Claim 3.2.2. *The solution is feasible, that is to say $z^*(C) \leq y_i^*$ for all $i \in \mathcal{M}$, $C \in \mathcal{C}_i$.*

It is easy to see that if the claims hold, then they also hold for a scaled solution $(\alpha \cdot y^*, \alpha \cdot z^*)$ with $\alpha > 0$. We can use this to reach an arbitrarily low objective value, thus proving that the dual is unbounded, which means the assumption in the beginning must be false. \square

Proof of Claim 3.2.1. The critical part of the prove is that $|\mathcal{M}(\mathcal{T}_{\text{BA}})| \leq |\mathcal{M}(\mathcal{T}_{\text{SA}})|$. Recall that moves of small jobs have a higher priority than moves of big jobs. This means that after a big-to-any blocker is added to the blocker tree, all available moves of small jobs will be added, before another big-to-any blocker can be. It is enough to show that whenever a big-to-any blocker is added, there is at least one small job j_S on the corresponding machine and there exists a machine $i' \in \Gamma(j_S) \setminus \{\sigma(j_S)\}$ that is not in $\mathcal{M}(\mathcal{T}_{\text{SA}} \cup \mathcal{T}_{\text{BA}})$. When this small job is moved, the big-to-any blocker will be deleted. This means as long as the big-to-any blocker exists, there will be a small-to-any blocker added after it, but before the next big-to-any blocker.

Let j_B be a big job and i some machine. We consider the situation where a big-to-any blocker for (j_B, i) is added. Assume w.l.o.g. that (j_B, i) is not valid. If it is valid, the blocker will be removed again instantly. At the time (j_B, i) is added as a big-to-any blocker, we have that $p(S_i(\mathcal{T}) + B_i) + p_{j_B} \leq 1 + R$. On the other hand, since (j_B, i) is not valid, we also know that $p(\sigma^{-1}(i)) + p_{j_B} > 1 + R$. This implies there has to be a small job j_S on i that is not in $S(\mathcal{T})$. By definition of $S(\mathcal{T})$ there must be a machine $i' \in \Gamma(j_S) \setminus \{i\}$ that is not in $\mathcal{M}(\mathcal{T}_{\text{SA}} \cup \mathcal{T}_{\text{BA}})$. Thus

$$\sum_{j \in \mathcal{J}} z_j^* \geq z_{j_{\text{new}}}^* + \sum_{i \in \mathcal{M}} z^*(\sigma^{-1}(i)) > 0 + \sum_{i \in \mathcal{M}} y_i^* + \frac{1}{6}|\mathcal{M}(\mathcal{T}_{\text{SA}})| - \frac{1}{6}|\mathcal{M}(\mathcal{T}_{\text{BA}})| \geq \sum_{i \in \mathcal{M}} y_i^*,$$

which concludes the proof. \square

Proof of Claim 3.2.2. Let $i \in \mathcal{M}$ and $C \in \mathcal{C}_i$. First, we consider the cases where $i \in \mathcal{M}(\mathcal{T}_{\text{SA}} \cup \mathcal{T}_{\text{BA}})$. Here we will show that $y_i^* \geq 1$, which implies the claim, since $z^*(C) \leq 1$.

Case 1. $i \in \mathcal{M}(\mathcal{T}_{\text{BA}})$. Then there must be a big-to-any move (j_B, i) which is not valid; that is to say, $p(\sigma^{-1}(i)) + p_{j_B} > 1 + R$. If every job $j \in \sigma^{-1}(i)$ has $p_j \leq 5/6$, then $p(\sigma^{-1}(i)) = z^*(\sigma^{-1}(i))$ and thus

$$y_i^* = p(\sigma^{-1}(i)) + \frac{1}{6} > 1 + R - p_{j_B} + \frac{1}{6} \geq 1.$$

Otherwise let $j'_B \in B_i$ with $p_{j'_B} > 5/6$. Then

$$y_i^* = z^*(\sigma^{-1}(i)) + \frac{1}{6} \geq z_{j'_B}^* + \frac{1}{6} = 1.$$

Case 2. $i \in \mathcal{M}(\mathcal{T}_{\text{SA}})$. Then there is a small move (j_S, i) , which is not valid, i.e., $p(\sigma^{-1}(i)) + p_{j_S} > 1 + R$. If there is more than one big job assigned to i , we easily see that $z^*(\sigma^{-1}(i)) \geq 1$. Otherwise at most one job can be rounded down, which implies $z^*(\sigma^{-1}(i)) \geq p(\sigma^{-1}(i)) - 1/6$. Hence

$$y_i^* \geq p(\sigma^{-1}(i)) - \frac{2}{6} > 1 + R - p_{j_S} - \frac{2}{6} \geq 1.$$

In the following, we study the cases where $i \notin \mathcal{M}(\mathcal{T}_{\text{SA}} \cup \mathcal{T}_{\text{BA}})$. Here it is sufficient to show that if there is a $j \in C$ big with $\sigma(j) \neq i$, then $z_j^* \leq z^*(\mathcal{A}_i(\mathcal{T}) \setminus C)$. The reason is that all small jobs in $C \cap \mathcal{A}(\mathcal{T})$ are in $\mathcal{A}_i(\mathcal{T})$ or else there would be a potential move. This means, if C does not contain a big job that are not also in $\sigma^{-1}(i)$, then

$$y_i^* = z^*(\sigma^{-1}(i)) \geq z^*(\mathcal{A}_i(\mathcal{T}) \cap \mathcal{J}_S) + z^*(B_i) \geq z^*(C).$$

Otherwise there is exactly one big $j_B \in C$ and $\sigma(j_B) \neq i$. Therefore

$$z^*(C) = z_{j_B}^* + z^*(\mathcal{A}_i(\mathcal{T}) \cap C) \leq z^*(\mathcal{A}_i(\mathcal{T}) \setminus C) + z^*(\mathcal{A}_i(\mathcal{T}) \cap C) = y_i^*.$$

For this purpose, in the cases below let $j \in C$ big with $\sigma(j) \neq i$.

Case 3. $i \in \mathcal{M}(\mathcal{T}_{\text{BB}}) \setminus \mathcal{M}(\mathcal{T}_{\text{SA}} \cup \mathcal{T}_{\text{BA}})$. We observe that Invariant 3 implies that $B_i \neq B_i^{\min}$ and thereby that there are at least two jobs j_B, j'_B in B_i . Thus

$$z_j^* \leq 5/6 < z_{j_B}^* + z_{j'_B}^* \leq z^*(\mathcal{A}_i(\mathcal{T}) \setminus C).$$

Case 4. $i \in \mathcal{M}(\mathcal{T}_{\text{BL}}) \setminus \mathcal{M}(\mathcal{T}_{\text{SA}} \cup \mathcal{T}_{\text{BA}} \cup \mathcal{T}_{\text{BB}})$. By Invariant 2 obviously B_i is not empty. Let $j_B = \min B_i$.

Case 4.1. $j > j_B$ and $p_{j_B} \leq 5/6$. We notice that j is not undesirable on i . If (j, i) is not a blocker in \mathcal{T} , then it must not be a potential move, that is

$$p(S_i(\mathcal{T}) \cup \{j_B\}) + p_j \geq p(S_i(\mathcal{T})) + p_j > 1 + R. \quad (1)$$

If it is a blocker, then by Invariant 2, inequality (1) holds as well, since the blocker must be a big-to-least blocker. Thus

$$z^*(\mathcal{A}_i(\mathcal{T}) \setminus C) \geq p(S_i(\mathcal{T}) \setminus C) + p_{j_B} \geq p(S_i(\mathcal{T})) - (1 - p_j) + p_{j_B} > R \geq z_j^*.$$

Case 4.2. $j < j_B$ or $p_{j_B} \geq 5/6$. Then $p_j \leq p_{j_B}$ or $z_{j_B}^* = 5/6 = z_j^*$; in particular

$$z_j^* \leq z_{j_B}^* \leq z^*(\mathcal{A}_i(\mathcal{T}) \setminus C).$$

Case 5. $i \notin \mathcal{M}(\mathcal{T})$. Then since (j, i) is not a potential move, we find $p(S_i(\mathcal{T})) + p_j > 1 + R$. Therefore

$$z_j^* \leq 1 + R - p_j - (1 - p_j) < p(S_i(\mathcal{T})) - (1 - p_j) \leq z^*(\mathcal{A}_i(\mathcal{T})) - z^*(\mathcal{A}_i(\mathcal{T}) \cap C) = z^*(\mathcal{A}_i(\mathcal{T}) \setminus C).$$

This completes the proof. \square

Theorem 3.3. *The algorithm terminates.*

Proof. We will define what we call the signature vector as

$$(\text{val}(\mathcal{B}_1), \text{val}(\mathcal{B}_2), \dots, \text{val}(\mathcal{B}_\ell), \infty),$$

where $\text{val}(\mathcal{B}_k)$ is the value of the move corresponding to a blocker \mathcal{B}_k at the time it was added to the tree. We argue that the signature vector decreases lexicographically after finitely many iterations of the algorithm's loop. Since we can bound the number of possible vectors fairly easily, we can also bound the overall running time: Each move can only appear once in the blocker tree, therefore $\ell \leq |\mathcal{J}| \cdot |\mathcal{M}|$. The range of val is also finite; hence the number of signature vectors is finite as well.

When the algorithm adds a blocker, clearly the signature vector decreases lexicographically. Now consider a valid move that is performed. It may be that many consecutive moves are performed, but at some point the algorithm will have to add a blocker again. We will observe only the last move (j_0, i_0) that is performed before a blocker is added. Let $\mathcal{B}_k = (j, i, \Theta)$ be the blocker that activated j_0 . In particular j_0 was assigned to i . We will show that the blocker that will be added in the next iteration has a lower value than \mathcal{B}_k had, thereby ultimately decreasing the lexicographic order. A candidate for this blocker is (j, i) . In the case distinction below we show that (j, i) is a potential move and $\text{val}(j, i)$ is lower than it was the last time the blocker was added. This means the next blocker that is added must definitely have a lower value. The blocker tree, that is $\mathcal{T}^{(\leq k-1)}$, is identical to the tree when (j, i) was added the last time. This means j is still not undesirable.

Case 1. **\mathcal{B}_k was a big-/small-to-any blocker.** Since all jobs were undesirable on i , no jobs have been moved there. If \mathcal{B}_k was a big-to-any blocker, then before j_0 was moved to i_0 , $p(S_i(\mathcal{T}^{(\leq k-1)}) \cup B_i) + p_j \leq 1 + R$ (Invariant 1) and this move cannot have increased the left-hand side. In other words, (j, i) is still a potential move of the same type. Since j_0 was removed from i , $|\sigma^{-1}(i)|$ has decreased and so has $\text{val}(j, i)$.

Case 2. **\mathcal{B}_k was a big-to-least blocker.** Because j_0 was activated by \mathcal{B}_k , it must have been the smallest big job on i . By Invariant 2 we know that before j_0 was moved,

$$p(S_i(\mathcal{T}^{(\leq k-1)})) + p_j \leq 1 + R \text{ and} \quad (2)$$

$$p(S_i(\mathcal{T}^{(\leq k-1)}) \cup B_i^{\min}) + p_j > 1 + R. \quad (3)$$

(2) still holds, since the move has no effect on it. If j_0 was the only big job, then $B_i = \emptyset$ and $p(S_i(\mathcal{T}) \cup B_i) + p_j \leq 1 + R$; (j, i) is now a big-to-any move, which has a lower value. If j_0 was not the only big job on i , then $\min B_i$ has increased, as no big jobs j_B with $j_B < j_0$ were moved to i since \mathcal{B}_k was added; the left-hand side of (3) has not decreased and (j, i) is still a big-to-least move, but has a lower value as well.

Case 3. **\mathcal{B}_k was a big-to-big blocker.** Since j_0 was undesirable regarding \mathcal{B}_k , it must be a big job. No big jobs were moved to i and j_0 was removed; hence $|B_i|$ has decreased. Before j_0 was moved,

$$p(S_i(\mathcal{T}^{(\leq k-1)})) + p_j \leq p(S_i(\mathcal{T}^{(\leq k-1)}) \cup B_i^{\min}) + p_j \leq 1 + R,$$

where we use Invariant 3 in the last inequality. The move obviously has no effect on the left-hand side and the inequality must still hold. It may, however, be that (j, i) is now a potential move of different type, but all other types have lower values; $\text{val}(j, i)$ decreases in any case. □

4 Conclusion

We have almost tripled the difference between integrality gap and the classical approximability result of 2. This also marks the gap between best estimation result and best approximation result (except for an arbitrary small $\epsilon > 0$).

On better bounds. The largest gap we have witnessed in an instance is $9/6$ (given in [7]). The bound on the integrality gap that was proven in this paper, $11/6$, is not necessarily tight.

An approach for improving the upper bound is to first consider a set of more constrained instances. A natural candidate for this could be the set of instances where for every job j either $p_j = 1$ or $p_j \leq 1/2$ (as before, 1 is the optimum of the configuration-LP). Even there it does not appear to be straight-forward to improve the algorithm, as we will show below. With the techniques by Svensson, it was already easy to see that these instances admit an integrality gap of no more than $11/6$. In this sense, we believe that a significant bottleneck that was introduced by the medium jobs in the original algorithm was eliminated by our improvement.

We will show that using certain basic assumptions on the proof we can construct a situation where the bound $11/6$ is the best we can get. In particular, the proof cannot be improved only by altering the constants. Let $1 + R$ be the makespan we try to achieve. First, we claim that we need to choose the value $z_{j_B}^*$ for big jobs j_B to be at most R : Let j_B be a big job with $p_{j_B} = 1$. We consider a situation where $i \in \Gamma(j_B)$ is a machine with $p(S_i(\mathcal{T})) = R + \epsilon$ for some $\epsilon > 0$ and no jobs besides $S_i(\mathcal{T})$ are assigned to i . The move (j_B, i) must not be a potential move, since Claim 3.2.1 relies on the existence of a small job on i that is not in $S(\mathcal{T})$ for every non-valid potential big-to-any move. Let C be the configuration with only j_B in it. Then for feasibility (Claim 3.2.2) we need that

$$z_{j_B}^* = z^*(C) \leq y_i^* = z^*(\sigma^{-1}(i)) = p(S_i(\mathcal{T})) = R + \epsilon.$$

Now let there be two machines $i \in \mathcal{M}(\mathcal{T}_{SA})$ and $i' \in \mathcal{M}(\mathcal{T}_{BA})$. i was added because of a move (j_S, i) of a small job j_S with $p_{j_S} = 1/2$ and it is assigned a big job j_B with $p_{j_B} = 1$. Then $p(\sigma^{-1}(i))$ can be as low as $1 + R - 1/2 + \epsilon$ for any $\epsilon > 0$ without (j_S, i) being valid. We assume that this is the case. Similarly we assume that $p(\sigma^{-1}(i')) = R + \epsilon$ and the big move to i' is not valid. In Claim 3.2.2 we rely on $y_i^* \geq 1$ and $y_{i'}^* \geq 1$ and in Claim 3.2.1 we need that $y_i^* + y_{i'}^* \leq z^*(\sigma^{-1}(i)) + z^*(\sigma^{-1}(i'))$. Together we get

$$2 \leq y_i^* + y_{i'}^* \leq z^*(\sigma^{-1}(i)) + z^*(\sigma^{-1}(i')) = \left(\frac{1}{2} + R + \epsilon - p_{j_B} + z_{j_B}^*\right) + (R + \epsilon) \leq 3R - \frac{1}{2} + 3\epsilon,$$

which is equivalent to $R \geq 5/6 - \epsilon$.

Beyond estimations. The perhaps most interesting question regarding this problem is, whether a polynomial approximation of the same guarantee (or at least of guarantee $2 - \delta$ for some constant $\delta > 0$) exists. Indeed, there are very few problems for which the best known estimation rate is lower than the best known approximation rate. Though it seems likely, we cannot know for sure that such an approximation algorithm exists. Feige and Jozeph even gave a proof that under some complexity assumptions there must be problems with estimation algorithms superior to the best approximation algorithms [5].

On the other hand, for the problem of RESTRICTED MAX-MIN FAIR ALLOCATION efficient variants of the respective local search algorithm were discovered [9, 1]. We conjecture that using similar approaches this is possible for the RESTRICTED ASSIGNMENT problem as well.

References

- [1] C. Annamalai, C. Kalaitzis, and O. Svensson. Combinatorial algorithm for restricted max-min fair allocation. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4-6, 2015*, pages 1357–1372, 2015.
- [2] N. Bansal and M. Sviridenko. The santa claus problem. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing, Seattle, WA, USA, May 21-23, 2006*, pages 31–40, 2006.
- [3] D. Chakrabarty, S. Khanna, and S. Li. On $(1, \epsilon)$ -restricted assignment makespan minimization. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4-6, 2015*, pages 1087–1101, 2015.
- [4] T. Ebenlendr, M. Krcál, and J. Sgall. Graph balancing: A special case of scheduling unrelated parallel machines. *Algorithmica*, 68(1):62–80, 2014.
- [5] U. Feige and S. Jozeph. Separation between estimation and approximation. In *Proceedings of the 2015 Conference on Innovations in Theoretical Computer Science, ITCS 2015, Rehovot, Israel, January 11-13, 2015*, pages 271–276, 2015.
- [6] C. Huang and S. Ott. A combinatorial approximation algorithm for graph balancing with light hyper edges. In *24th Annual European Symposium on Algorithms, ESA 2016, August 22-24, 2016, Aarhus, Denmark*, pages 49:1–49:15, 2016.
- [7] K. Jansen, K. Land, and M. Maack. Estimating the makespan of the two-valued restricted assignment problem. In *15th Scandinavian Symposium and Workshops on Algorithm Theory, SWAT 2016, June 22-24, 2016, Reykjavik, Iceland*, pages 24:1–24:13, 2016.
- [8] J. K. Lenstra, D. B. Shmoys, and E. Tardos. Approximation algorithms for scheduling unrelated parallel machines. *Mathematical Programming*, 46(3):259–271, 1990.
- [9] L. Poláček and O. Svensson. Quasi-polynomial local search for restricted max-min fair allocation. *ACM Transactions on Algorithms*, 12(2):13, 2016.
- [10] P. Schuurman and G. J. Woeginger. Polynomial time approximation algorithms for machine scheduling: ten open problems. *Journal of Scheduling*, 2(5):203–213, 1999.
- [11] O. Svensson. Santa claus schedules jobs on unrelated machines. *SIAM Journal on Computing*, 41(5):1318–1341, 2012.
- [12] J. Verschae and A. Wiese. On the configuration-lp for scheduling on unrelated machines. *Journal of Scheduling*, 17(4):371–383, 2014.
- [13] D. P. Williamson and D. B. Shmoys. *The Design of Approximation Algorithms*. Cambridge University Press, 2011.